

Attorney Docket No.  
212160

**PATENT APPLICATION**

Invention Title:

NON-INVASIVE LATENCY MONITORING IN A STORE-AND-FORWARD  
REPLICATION SYSTEM

Inventor:

Jeffrey B. Parham	U.S.	Redmond	Washington
INVENTOR'S NAME	CITIZENSHIP	CITY OF RESIDENCE	STATE or FOREIGN COUNTRY

Be it known that the inventor listed above has invented a certain new and useful invention with the title shown above of which the following is a specification.

# NON-INVASIVE LATENCY MONITORING IN A STORE-AND-FORWARD REPLICATION SYSTEM

## AREA OF THE INVENTION

5           The invention relates generally to the transmission of data to multiple computers in a computer network and, more particularly, to a method of monitoring data updates in complex replicated systems.

## BACKGROUND OF THE INVENTION

10           In the field of computer networking, many efforts have been made to develop the most efficient and reliable way for servers within a computer network to communicate updates with one another. In particular, the problem of monitoring the status of servers and server updates has been a challenge given that most conventional network database systems often contain multiple servers that are geographically dispersed. "Updates" are generally  
15           modifications to data objects or attributes of data objects within the distribute database. Because copies of the data objects may exist on several servers, any modification made to an object of a database on one of the servers must be propagated to all other copies of the object in the system such that updated objects at the various servers reflect the modification.

20           In order to keep the data objects maintained by each server current, updates made by one server are "replicated" to all other servers of the database system by a process called "replication." During replication, a "source" server sends data to a "destination" server, and the destination server updates its database with the objects that were modified. The updates being replicated may have originated on the source server or on another server in the system. A server on which a modification to an object is initially made, rather than an update  
25           received through a replication, is referred to as an "originating" server. Ultimately, the goal is to replicate an update message to all servers in the system that require updated object information.

30           Several techniques have been developed to effectuate replication in a network. One such technique is applicable in a simple replicated system, whereby a small number of servers holding the replicated data, referred to as "replicas" or "replica servers," can be directly connected to all servers (or, as an optimization, all read/write (or master) replicas). This

technique, referred to as "full-mesh replication," operates in a manner such that as an update is originated on one replica, that replica then sends the update directly to every other replica. According to this technique, if there are  $N$  replicas, each replica in a full-mesh system has  $N-1$  inbound communication paths and  $N-1$  outbound communication paths. The network that connects the replicas must sustain all of the order  $N^2$  communication paths. At some value  $N$ , however, the load on the individual replica servers and/or on the network becomes too great to effectively sustain full-mesh replication.

Consequently, reducing the load requires a reduction in the number of overall communication paths between servers. A workable solution calls for a model that does not require direct connections between all replica servers. One technique, referred to as "store-and-forward replication," overcomes the loading problem associated with full-mesh replication by allowing an update to originate on server A, replicate to server B, and then replicate from server B to server C. This method, wherein an update can flow through intermediary replicas, does not require a connection between destination server C and originating server A.

Compounding the problem of replication, however, is the fact that replica servers are prone to failure. Failures can be due to a number of problems including network errors, hardware problems, server or network misconfiguration, etc. Some failures result from the situation where the replica server is simply not functioning. In that instance, other servers that query the non-functioning server for information will realize that the server is not operational. Other failures, however, result from the situation where the replica server is functioning but not receiving updates originating from other replica servers. In this more troubling situation, clients that contact a replica server in such a state receive increasingly outdated data until the failure is fixed or the replica server is taken offline. In the meantime, the effects of the outdated data can be subtle. For example, an address book might report an old phone number for a contact. These subtle effects might go unnoticed for a long period of time, during which the amount of data that needs to be replicated to the failed replica server continues to grow as does the amount of time it will take to bring the replica server up-to-date once the failure has been identified and corrected. Thus, most replica systems provide some

form of service feature to monitor replica servers so that failures preventing replication can be identified and thereafter rectified.

Monitoring the replication state of a replica server in a full-mesh system is typically easy. Because each server communicates with all other servers directly, each replica server maintains state information about which updates it has successfully sent or received to or from each other replica server. This information, referred to as the "direct replica partner state," contains timestamps used to evaluate the integrity of the data. For example, the timestamp represents the time of the last successful replication or the time of the last replication attempt. The direct replica partner state, thus, can be queried by an administrator or monitoring tool to determine if replication is functioning properly. Unfortunately, replication monitoring of a full-mesh system becomes an impractical solution for the reason that directly querying all servers in a system is impractical.

Monitoring the replication state of a replica server in a store-and-forward replication system, however, is more difficult. The direct replica partner state, which by definition does not include information about replica servers that are not direct partners, yields only a partial view of the quality of the data replicated inbound or outbound from a given replica server. For example, if server A replicates to/from only server B and server B replicates to/from only server C, then examining server C alone would not provide the administrator with any information regarding how current the data on server C is with respect to the data on server A.

Another method for monitoring the health of store-and-forward systems, employed by Microsoft's Operations Manager, requires each replica server to originate an artificial update referred to as a "heartbeat." This update is made solely for the benefit of monitoring the replication system. If all replicas are originating artificial updates – which is as simple as writing the current time on data related to that replica server – then any given replica server can examine its local copy of the data, which includes the data being written by all other replica servers, to determine the last update it received from every other replica server. If the last update it received originated at some time in the past beyond what is expected, then action is taken to inform an administrator to investigate the server failure. One problem with this self-

monitoring approach is that it requires the data corresponding to each replica server to be updated periodically, often with a period less than the maximum latency tolerated between replica servers, in order to be monitored. As such, the origination and replication of these write messages can be very expensive.

5 As the number of replica servers in the replication system grows, the complexity of determining how current any given replica server is with respect to all other replica servers grows very quickly, both in terms of the number of servers that must be queried and in determining which servers must be queried. In view of the foregoing, it can be seen that there is a need for a method for proactively monitoring replica servers to ensure that failures preventing replication are addressed and rectified in a timely and efficient manner.

## SUMMARY OF THE INVENTION

The present invention is directed to a method for monitoring replica servers to ensure that failures preventing replication of object updates are identified and rectified in a timely and efficient manner.

15 According to aspects of the present invention, a networked computer system comprises a plurality of servers connected by data links. Each server in the networked computer system periodically replicates object updates from another server in the system during replication. Associated with each server in the networked computer system is a replica partner vector table that includes relevant information about all servers in the system. In particular, each server maintains an independent, monotonically increasing update sequence number (USN) that it associates with each update it originates. The (originating server ID, USN) pair associated with each update is replicated along with the updated data. The replica partner vector table is made up of a number of (originating server ID, USN) pairs. Each such pair is an assertion that this server has applied all updates originated by the server corresponding to the originating server ID at or before the given USN. These assertions are communicated to other servers during replication to avoid replicating updates to a server if the server has already received and applied them. As such, other servers in the system use the pair to identify the current state of their own object data.

The replica partner vector table also includes timestamp information that identifies the time of its last update and the last successful replication attempt for each replica server in the system, regardless of whether the replica communicated an update to that replica directly. Thus, each replica maintains information about its direct and indirect partners such that the state of indirect partners flows through the direct partners.

During the replication process, a server transmits its replica partner vector table to the replica server. The replica server then compares its own replica partner vector table with the replica partner vector table it receives from the server. By comparing the USN and the timestamp, the replica server can determine what updates the server is missing. Thereafter, the replica server transmits the missing object updates in addition to the new USNs and timestamps. After receiving the object updates, the server updates its own replica partner vector table to include the new USNs and timestamps.

A monitoring tool running at each server, separate from the replication process, evaluates the USN and timestamp entries in the replica partner vector table. As part of the monitoring process, all timestamps are compared to the current time to determine when the last successful replication attempt occurred. If the time of the last successful replication attempt as compared to the current time exceeds an acceptable latency period, an alert is generated whereby users and/or a network administrator are informed of the problem.

Additional features and advantages of the invention will be made apparent from the following detailed description of illustrative embodiments that proceeds with reference to the accompanying figures.

## BRIEF DESCRIPTION OF THE DRAWINGS

While the appended claims set forth the features of the present invention with particularity, the invention, together with its objects and advantages, may be best understood from the following detailed description taken in conjunction with the accompanying drawings of which:

FIG. 1 is an example of a networked computer system in which aspects of the present invention and/or portions thereof may be incorporated;

FIG. 2 is a block diagram of a general purpose computer in which aspects of the present invention and/or portions thereof may be incorporated;

FIG. 3 depicts an exemplary networked computer system with several servers implementing store-and-forward replication with one replication path;

FIG. 4 depicts an exemplary networked computer system with several servers implementing store-and-forward replication with multiple replication paths;

FIGS. 5a-5b depict exemplary replica partner vector tables, illustrating updates thereto, in accordance with the present invention;

FIGS. 6a-6d show a simple store-and-forward replication system with several servers, each with an associated replica partner vector table, illustrating updates to the replica partner vector tables in accordance with one embodiment of the present invention; and

FIG. 7 is a flow diagram of an exemplary replication monitoring process in accordance with one embodiment of the present invention.

## DETAILED DESCRIPTION OF ILLUSTRATIVE EMBODIMENTS

Although it is not required, the present invention may be implemented by program modules that are executed by a computer. Generally, program modules include routines, objects, components, data structures and the like that perform particular tasks or implement particular abstract data types. A program may include one or more program modules. The invention may be implemented on a variety of types of computers, including personal computers (PCs), hand-held devices, multi-processor systems, microprocessor-based programmable consumer electronics, network PCs, minicomputers, mainframe computers and the like. The invention may also be employed in distributed computing environments, where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, modules may be located in both local and remote memory storage devices.

An example of a networked environment in which this system may be used will now be described with reference to FIG. 1. The example network includes several computers 100a-f communicating with one another over a network 102, represented as a cloud. Network

102 may include any of many well-known components, such as routers, gateways, hubs, etc. and may allow computers 100a-f to communicate via wired and/or wireless media.

Referring to **FIG. 2**, an example of a basic configuration for a computer on which the system described herein may be implemented is shown. In its most basic configuration, computers 100a-f typically include at least one processing unit 112 and memory 114. Depending on the exact configuration and type of the computer, the memory 114 may be volatile (such as RAM), non-volatile (such as ROM or flash memory) or some combination of the two. This most basic configuration is illustrated in **FIG. 2** by dashed line 106. Additionally, the computer may also have additional features/functionality. For example, computers 100a-f may also include additional storage (removable and/or non-removable) including, but not limited to, magnetic or optical disks or tape. Computer storage media includes volatile and non-volatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program modules, or other data. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disk (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to stored the desired information and which can be accessed by computers 100a-f. Any such computer storage media may be part of computers 100a-f.

Computers 100a-f may also contain communications connections that allow the device to communicate with other devices. A communication connection is an example of a communication medium. Communication media typically embodies computer readable instructions, data structures, program modules or other data in a modulated data signal such as a carrier wave or other transport mechanism and includes any information delivery media. By way of example, and not limitation, communication media includes wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, RF, infrared and other wireless media. The term computer readable media as used herein includes both storage media and communication media.



Computers 100a-f may also have input devices such as a keyboard, mouse, pen, voice input device, touch input device, etc. Output devices such as a display 116, speakers, a printer, etc. may also be included. All these devices are well known in the art and need not be discussed at length here.

**FIG. 3** depicts an exemplary networked computer system with several servers. As shown in **FIG. 3**, the computer system 120 includes a plurality of servers (referenced as  $S_1$ ,  $S_2$ ,  $S_3$ ,  $S_4$ , and  $S_5$ ) interconnected by data links 122. The servers may be any appropriate server, for example, a central server or a client server. Any server is employed in various embodiments of the present invention. The data links 122 comprise any appropriate data link, for example, a local area network, a wide area network, or the Internet. Various data links are employed in alternative embodiments of the invention. Servers and data links are generally known to the relevant public and therefore need not be described herein in any detail. Furthermore, as is known in the industry, conventional networked computer systems contain large numbers of distributed servers interconnected by varied data links.

As shown in **FIG. 3**, the computer system 120 implements store-and-forward replication with one replication path 124. As is typical in a store-and-forward replication system, each replica server is directly connected to a number of other replica servers such that each server receives "direct updates" from its "direct partners." For example,  $S_3$  receives direct updates from its two direct partners,  $S_2$  and  $S_4$ . Similarly, each server receives "indirect updates" from each "indirect partner" in the replication system. This method allows indirect updates from a "source server" to flow through direct partners and arrive at the "destination server." For example, server  $S_5$  receives indirect updates from indirect partner  $S_1$  by way of replica servers  $S_2$ ,  $S_3$  and  $S_4$ . In other words, an update message is sent along path 124 from server  $S_1$  to server  $S_5$ . Of course, if any server along path 124 is not functioning, the update message will not arrive at destination server  $S_5$ .

To avoid this problem, conventional store-and-forward systems include multiple replication paths. **FIG. 4** depicts a simple store-and forward replication system 126 with two replication paths 128, 130. In this example, an update message forwarded by source server  $S_1$  to destination server  $S_5$  takes one of two paths. According to one path 128, the update message

from source server  $S_1$  is sent through indirect partners  $S_2$  and  $S_3$ , serially, and then on to direct partner  $S_4$  which finally forwards the update to destination server  $S_5$ . The other path 130 requires that the update message generated by source server  $S_1$  be sent to  $S_6$ . Server  $S_6$ , a direct partner of  $S_5$ , then sends the update on to destination server  $S_5$ . Because the replication system

5 contains multiple replication paths, a single, non-functioning server does not prevent an update message from arriving at the destination server.

Of course, in order for this replication method to operate effectively, each replica server in a replication system contains information about both its direct partners and indirect partners. This information kept by the replica server is referred to as the "replica partner state." According to an exemplary embodiment of the invention, replica servers represent the replica state for all replica server partners, both direct and indirect, as a "replica partner vector." Each entry in the replica partner vector table corresponds to exactly one replica server such that the replica partner vector represents the entirety of a given replica server's knowledge of its replication state with respect to all other replica servers in the

10 networked computer system.

According to one aspect of the exemplary embodiment of the invention, replica servers in the replication system have a replica partner vector table associated with it. Depicted in **FIG. 5a** is an exemplary replica partner vector table 132 representing the replica partner state for server  $S_3$  shown in **FIG. 4**. The replica partner vector 132 includes a data field 134

15 corresponding to each replica server in the replication system 126. For example, replica partner vector table 132 includes references to  $S_1$ ,  $S_2$ ,  $S_3$ ,  $S_4$ ,  $S_5$ , and  $S_6$ . Alternatively, server  $S_3$  includes references in its vector table to all servers except the reference relating to its own state, which is then stored separately.

Each replica server representation in the replica partner vector table 132

25 includes a data field 136 representing an update sequence number (USN). According to an embodiment of the invention, the USN is a serial number assigned sequentially to each update recorded by the replica server. For example, as shown in **FIG. 5a**, the replica partner state of  $S_3$  indicates that  $S_3$  is current up to a USN value 37 for  $S_1$ , a USN value 202 for  $S_2$ , a USN value 64 for  $S_4$ , a USN value 189 for  $S_5$ , and a USN value 101 for  $S_6$ . An implementation

example of such a USN is found in the MICROSOFT ACTIVE DIRECTORY service (a product of Microsoft Corp. of Redmond, Washington).

Each replica server representation in the replica partner vector table 132 includes a data field 138 representing a timestamp referred to as the "last successful replication timestamp" (LSRT). The LSRT for any replica server specifies the last time at which a successful replication attempt was made by that replica server. According to one aspect of the invention, at the end of a successful replication attempt, the LSRT will be posted regardless of whether any updates were replicated from server to server. Importantly, for purposes of monitoring servers within a networked computer system, this method provides a way to differentiate between a functioning server with no updates and a non-functioning server incapable of replicating its update information.

Thus, USN data field 134 and LSRT data field 138, together, provide information on both the status of the server and objects updates at the server. For example, as demonstrated in replica partner vector table 132 shown in **FIG. 5a**, server S<sub>3</sub> replicated object updates for all servers in the computer system including server S<sub>4</sub>. In particular, data field 136 indicates that S<sub>3</sub> is current up to a USN value 64 for server S<sub>4</sub> and data field 138 indicates that the last successful replication attempt for data relating to server S<sub>4</sub> occurred on 11/13/2001 at 06:03 a.m. The LSRT may be any appropriate absolute representation of time capable of comparison, for example, 2001/11/13 or 11132001. What is known from the replica partner state data at server S<sub>3</sub> is that server S<sub>3</sub> received updates relating to server S<sub>4</sub> from another server in the computer system 126 either directly from server S<sub>4</sub> or indirectly from another server. What is not known is whether a USN value 64 corresponds to the most up-to-date information at server S<sub>4</sub>. As is typical in a replication system, server S<sub>3</sub> will continue to replicate updates with its replica partners such that it will receive any updates subsequent to a USN value 64 from server S<sub>4</sub>.

For example, if server S<sub>4</sub> posts an update having a USN value equal to 65 and server S<sub>3</sub> replicates the update, then the replica partner vector table 132 of server S<sub>3</sub> will reflect the new replica partner state. **FIG. 5b** depicts server S<sub>3</sub>'s replica partner vector table 140 with data field 142 corresponding to each replica server in the replication system 126, data field 144

representing the USN of each update, and data field 146 representing the LSRT. **FIG. 5b** shows the replica partner state after an update having a USN value 65 is replicated to server S<sub>3</sub>. In particular, data field 144 indicates that S<sub>3</sub> is current up to a USN value 65 for server S<sub>4</sub> and data field 146 indicates that the last successful replication attempt for data relating to server S<sub>4</sub> occurred on 11/13/2001 at 10:29 a.m. Also shown in **FIG. 5b**, servers S<sub>1</sub>, S<sub>2</sub>, S<sub>5</sub> and S<sub>6</sub> did not replicate object updates to S<sub>3</sub> because S<sub>3</sub> was already up-to-date with all object updates as indicated by the USN and LSRT.

According to another aspect of an exemplary embodiment of the invention, the server requesting replication transmits its replica partner vector table to the other server. The server requesting replication is referred to as the "local" server and the server being queried is referred to as the "remote" server. To facilitate replication, the remote server, upon receiving the replica partner vector table from the local server, compares the replica partner vector table from the local server to its own replica partner vector table. By comparing the USN and LSRT for each entry in the replica partner vector table, the remote server determines whether any object updates need to be replicated to the local replica server. Thereafter, any object updates or timestamp updates are replicated (i.e., transmitted) to the local server, after which, the local server improves its own replica partner vector table to reflect the replication updates.

For example, **FIG. 6a** depicts an exemplary networked computer system 150 which includes a plurality of servers (referenced as S<sub>1</sub>, S<sub>2</sub>, and S<sub>3</sub>) interconnected by data links 152. Servers S<sub>1</sub>, S<sub>2</sub>, and S<sub>3</sub> are further connected to other servers in the network 154 via data links 152. Replica partner vector tables 156, 158, and 160 include a USN data field and an LSRT data field, for each server S<sub>1</sub>, S<sub>2</sub>, and S<sub>3</sub> in the networked computer system 150, to facilitate replication and status monitoring of servers and server updates.

As shown in **FIG. 6a**, replica partner vector table 156 for server S<sub>1</sub> indicates that the last object update to server S<sub>1</sub> was sequentially assigned a USN value 42 and successfully updated at 11:02 a.m. on 12/07/01, as indicated by the LSRT data field. Replica partner vector table 156 for server S<sub>1</sub> also indicates that server S<sub>1</sub> replicated an object update having a USN value 101 from server S<sub>2</sub> and that the last successful replication attempt for server S<sub>2</sub> occurred at 9:02 a.m. on 12/07/01, as indicated by the LSRT data field. Replica partner vector table 156

for server S<sub>1</sub> further indicates that server S<sub>1</sub> replicated an object update having a USN value 67 from S<sub>3</sub> and that the last successful replication attempt for server S<sub>3</sub> occurred at 11:46 p.m. on 12/06/01.

As shown in **FIG. 6a**, replica partner vector table 158 for server S<sub>2</sub> indicates that the last object update to server S<sub>2</sub> was sequentially assigned a USN value 102 and successfully updated at 11:01 a.m. on 12/07/01, as indicated by the LSRT data field. Replica partner vector table 158 for server S<sub>2</sub> also indicates that server S<sub>2</sub> replicated an object update having a USN value 41 from server S<sub>1</sub> and that the last successful replication attempt from server S<sub>1</sub> occurred at 8:31 a.m. on 12/07/01, as indicated by the LSRT data field. Replica partner vector table 158 for server S<sub>2</sub> further indicates that server S<sub>2</sub> replicated an object update having a USN value 67 from S<sub>3</sub> and that the last successful replication attempt for server S<sub>3</sub> occurred at 11:46 p.m. on 12/06/01.

As shown in **FIG. 6a**, replica partner vector table 160 for server S<sub>3</sub> indicates that the last object update to server S<sub>3</sub> was sequentially assigned a USN value 67 and successfully updated at 11:03 a.m. on 12/07/01, as indicated by the LSRT data field. Replica partner vector table 160 for server S<sub>3</sub> also indicates that server S<sub>3</sub> replicated an object update having a USN value 42 from server S<sub>1</sub> and that the last successful replication attempt for server S<sub>1</sub> occurred at 11:02 a.m. on 12/07/01, as indicated by the LSRT data field. Replica partner vector table 160 for server S<sub>3</sub> further indicates server S<sub>3</sub> replicated an object update having a USN value 102 from server S<sub>2</sub> and that the last successful replication attempt for server S<sub>3</sub> occurred at 11:01 a.m. on 12/07/01.

According to one embodiment of the invention, as server S<sub>1</sub> attempts replication from server S<sub>2</sub>, server S<sub>1</sub> transmits its replica partner vector table 156 to server S<sub>2</sub>. In another embodiment of the invention, server S<sub>1</sub> and server S<sub>2</sub> exchange their respective replica partner vector tables 156, 158. During replication, server S<sub>2</sub>, upon receiving the replica partner vector table 156, compares the replica partner vector table 156 to its own replica partner vector table 158. A comparison of all entries in the USN data field and LSRT data field indicates that the data of server S<sub>2</sub> at server S<sub>1</sub> (i.e., USN value 101, LSRT value 12/07/01, 9:02 a.m.) is not current with respect to data at server S<sub>2</sub> (i.e., USN value 102, LSRT value 12/07/01, 11:01

a.m.). This comparison also indicates that the data of server  $S_1$  at server  $S_2$  (i.e., USN value 41, LSRT value 12/07/01, 8:31 a.m.) is not current with respect to data at server  $S_1$  (i.e., USN value 42, LSRT value 12/07/01, 11:02 a.m.). During the replication process, server  $S_1$  also compares the data of server  $S_3$  at server  $S_2$  (i.e., USN value 67, LSRT value 12/06/01, 11:46 p.m.) with its own data relating to server  $S_3$  (i.e., USN value 67, LSRT value 12/06/01, 11:46 p.m.) revealing that the replica servers have the same updates and the same timestamps. After comparison of all other server data is complete, servers  $S_1$  and  $S_2$  replicate (i.e., transmit) any necessary updates and associated timestamps.

**FIG. 6b** depicts replica partner vector tables 156, 158 after server  $S_2$  completes replication of the object data from server  $S_1$  and server  $S_1$  completes replication of the object data from server  $S_2$ . Server  $S_1$  has improved its own replica partner vector table 156 to reflect a USN value 102 and an LSRT value 12/07/01, 11:01 a.m. indicating a successful replication attempt. Similarly, server  $S_2$  has improved its own replica partner vector table 158 to reflect a USN value 42 and an LSRT value 12/07/01, 11:02 a.m. also indicating a successful replication attempt.

As shown in **FIG. 6b**, as server  $S_2$  then attempts replication from server  $S_3$ , server  $S_2$  transmits its improved replica partner vector table 158 to server  $S_3$  and vice-versa. Server  $S_3$ , upon receiving the replica partner vector table 158, compares the replica partner vector table 158 to its own replica partner vector table 160. A comparison of all entries in the USN data field and LSRT data field indicates that the data related to servers  $S_1$  and  $S_2$  are updated with respect to the data in replica partner vector table 160, however, the data relating to server  $S_3$  at server  $S_2$  (i.e., USN value 67, LSRT value 12/06/01, 11:46 p.m.) is not current with respect to data at server  $S_3$  (i.e., USN value 67, LSRT value 12/07/01, 11:03 a.m.). In situations where the USN is the same (e.g., both replica servers have the same set of updates), but the LSRT is different, the local server, in this case server  $S_2$ , improves its timestamp to be the same as the remote server, server  $S_3$ . Thus, **FIG. 6c** depicts replica partner vector table 158 after server  $S_2$  replicates the timestamp data from server  $S_3$ .

As a monitoring tool, the timestamp having a value of 11:03 a.m. indicates that server  $S_3$  is functioning as opposed to being idle in the best-case scenario or non-functioning in

the worst-case scenario. **FIG. 7** is a flow diagram of an exemplary replication monitoring process performed by a replica server  $S_1$  in the system. As shown in **FIG. 7**, at step 200, a counter  $N$  is set to 1. At step 202, the variable  $T_{\max\_latency}$  is set to the maximum allowable latency between two replica servers. In this example,  $T_{\max\_latency}$  is set to 60 minutes. However,

5 it should be appreciated that the maximum allowable latency between replica servers in a network computer system can be set to any value, but is optimized, for example, for the best possible performance of the system. Optimization factors include the total time it takes a server to reset after a temporary failure and the acceptable dormancy of inactive servers.

At step 204,  $T_{\text{now}}$  is set to the current time within the network computer system.

10 Next, at step 206, the monitoring method determines whether the counter  $N$  is greater than the total number of servers in the system. If the counter  $N$  is not greater, then at step 208 the variable  $T_{\text{last}}$  is set to the LSRT of  $S_N$ .  $T_{\text{last}}$ , therefore, represents the time at which the last successful replication was performed between the replica server  $S_1$  and server  $S_N$ . At step 210, the monitoring method calculates the latency,  $T_{\text{now}}$  minus  $T_{\text{last}}$ , and determines whether the latency value is greater than the maximum allowable latency,  $T_{\max\_latency}$ .

15 If the latency is greater, then at step 212 an alert is generated. According to the invention, an alert may be any action that informs the user or network administrator that server  $S_N$  is experiencing problems. By way of example, and not limitation, possible alerts include sending an email to users of the network system, displaying a message on the user's computer screen, sending a message or email to a network administrator or central monitoring, or attempting to repair the problematic server. After the alert is generated, or in the case where the latency is not greater, the monitoring method proceeds to step 214 where counter  $N$  is increased by one so that the next server can be examined. Returning to step 206, the monitoring method again determines whether the counter  $N$  is greater than the total number of

20 servers in the system.

25

When counter  $N$  exceeds the number of servers in the network system indicating that the monitoring method has completed latency evaluation of all the servers identified in the replica partner vector table for replica server  $S_1$ , the monitoring method ends at step 216.

Alternatively, the monitoring method may proceed to step 200 where counter N is reset to 1, thus, allowing the monitoring process to repeat indefinitely.

In the example shown in **FIG. 6d**, server S<sub>4</sub> is interconnected to the networked computer system 150 via data link 152. Server S<sub>4</sub> includes a replica partner vector table 162 that indicates that the last object update to server S<sub>4</sub> was assigned a USN value 31 and successfully updated at 4:15 p.m. on 12/06/01. Replica partner vector table 162 also indicates server S<sub>4</sub> replicated object updates from server S<sub>1</sub> (i.e., USN value 22, LSRT value 12/06/01, 4:01 p.m.), server S<sub>2</sub> (i.e., USN value 71, LSRT value 12/06/01, 4:07 p.m.), and server S<sub>3</sub> (i.e., USN value 56, LSRT value 12/06/01, 3:59 p.m.).

By way of example, the monitoring method described in **FIG. 7** is explained with respect to the system shown in **FIG. 6d**, in particular server S<sub>2</sub>. At steps 200-204, counter N is set to 1, T<sub>max\_latency</sub> is set to 60 minutes, and T<sub>now</sub> is set to 12/07/01, 11:05 a.m. At step 206, counter N is compared to the total number of servers in networked computer system 150, a number greater than four because servers S<sub>1</sub>, S<sub>2</sub>, S<sub>3</sub>, and S<sub>4</sub> are connected to other servers in the network 154. At step 208, T<sub>last</sub> is set to S<sub>1</sub>(LSRT) which according to replica partner vector table 158 is 12/07/01, 11:02 a.m. A comparison at step 210 reveals that T<sub>now</sub> minus T<sub>last</sub> (i.e., 12/07/01, 11:05 a.m. minus 12/07/01, 11:02 a.m., or 3 minutes) is not greater than maximum allowable latency, T<sub>max\_latency</sub>, of 60 minutes. As a result, the monitoring method proceeds to step 214, where counter N is incremented to 2, and then on to step 206.

At step 206, counter N is again compared to the total number of servers in networked computer system 150. Because N is less, at step 208, T<sub>last</sub> is set to S<sub>2</sub>(LSRT) which according to replica partner vector table 158 is 12/07/01, 11:01 a.m. A comparison at step 210 reveals that T<sub>now</sub> minus T<sub>last</sub> (i.e., 12/07/01, 11:05 a.m. minus 12/07/01, 11:01 a.m., or 4 minutes) is not greater than maximum allowable latency, T<sub>max\_latency</sub>, of 60 minutes. Importantly, since the monitoring method is running on server S<sub>2</sub>, an evaluation of the latency of S<sub>2</sub>(LSRT) demonstrates an embodiment of the invention whereby a server can conduct self-monitoring of its own updates. Since a server is constantly replicating its own updates and the LSRT is updated to reflect the current time, the latency in this situation should be zero or near zero if the server is functioning properly.



Counter N is then incremented to 3 at step 214 and another comparison is made at step 206. At step 208,  $T_{last}$  is set to  $S_3$ (LSRT) which according to replica partner vector table 158 is 12/07/01, 11:03 a.m. A comparison at step 210 reveals that  $T_{now}$  minus  $T_{last}$  (i.e., 12/07/01, 11:05 a.m. minus 12/07/01, 11:03 a.m., or 2 minutes) is not greater than maximum allowable latency,  $T_{max\_latency}$ , of 60 minutes.

After counter N is incremented to 4 at step 214 and another comparison is made at step 206. At step 208,  $T_{last}$  is set to  $S_4$ (LSRT) which according to replica partner vector table 158 is 12/06/01, 4:15 p.m. A comparison at step 210 reveals that  $T_{now}$  minus  $T_{last}$  (i.e., 12/07/01, 11:05 a.m. minus 12/06/01, 4:15 p.m., or 18 hours, 10 minutes) is greater than maximum allowable latency,  $T_{max\_latency}$ , of 60 minutes and thus, an alert is generated at step 212. After the alert is generated, counter N is once again incremented and the method continues until all servers in the computer system 150 have been evaluated.

It can thus be seen that a new and useful method for monitoring updates to replica servers in a replication system has been provided. In view of the many possible embodiments to which the principles of this invention may be applied, it should be recognized that the embodiments described herein with respect to the drawing figures is meant to be illustrative only and should not be taken as limiting the scope of invention. For example, those of skill in the art will recognize that the elements of the illustrated embodiments shown in software may be implemented in hardware and vice versa or that the illustrated embodiments can be modified in arrangement and detail without departing from the spirit of the invention. Therefore, the invention as described herein contemplates all such embodiments as may come within the scope of the following claims and equivalents thereof.